CREATING AND INHABITING A SPACE: REAL-TIME ALGORITHMS AND PERFORMANCE ISSUES IN ROGER REYNOLDS'S *ANGELSPACE* (2008-2009)

Jacob David Sudol
University of California, San Diego
Department of Music

Roger Reynolds
University of California, San Diego
Department of Music

ABSTRACT

The authors present the development and current performance practice of Roger Reynolds real-time interactive computer music composition *AngelSpace* (2008-2009)[1]. The essay first gives a broad introduction to the composition's inspiration, materials, and form. The authors then describe the development and details of the spatialization and the other deterministic and probabilistic real-time algorithms utilized in each compositional section or "stage" of *Angel Space*. These descriptions also address performance strategies that the authors have discovered during experiments and rehearsals. The conclusion reflects on the aesthetic and technical directions suggested by the new compositional approach explored in *AngelSpace.*

1. INTRODUCTION

1.1. Background

AngelSpace [4] is the second work in Roger Reynolds's series of extended real-time interactive computer music compositions. Reynolds created the first work in this series, *The Image Machine* (2005) [3], with his musical assistant at the time Pei Xiang. This previous work uses pre-recorded sound files that Reynolds created for the collaborative work *22* (2004-2005). In *The Image Machine* a performer operates a Max/MSP patch to unfold a predetermined form and control a number of algorithms that reorder, overlap, and process these files in real-time.

1.2. Materials and Form in *AngelSpace*

AngelSpace arose from Reynolds's wish to create a work similar to *The Image Machine* with materials from his chamber piano concerto *The Angel of Death* (1998-2001) [2] that was commissioned by IRCAM. The latter composition arises out of a set of five thematic elements that are performed in closely parallel versions by either a piano soloist or a chamber orchestra. This thematic material is extrapolated in two formal sections, each with its own aesthetic imperative. One part is referred to as Sectional (S), the other as Domain (D). These distinctions have to do with whether the materials are formally treated as delimited, defined, and marked by clear boundaries (S) or rather organically flowing and morphing from one moment to the next (D).

In addition to their use in the instrumental music, the same thematic elements (T1 – T5) were recorded and then transformed into ten six-channel computer "images" (D1, D2, D5, D6, D9, D10 and S3, S4, S7, S8). *AngelSpace* utilizes this array of materials – Thematic elements and Computer images – as the material that defines a broadly conceived, environment or semi-improvisatory "space" within which assertion, expectation, insistence, musings, and more, can be explored by a performer in widely varying weave.

AngelSpace presents the material in eleven sections or "stages" (see Figure 1). The work – in its conception, programmed realization, and performance – involves constant assessment regarding the balance between continuity (an ongoing feeling) and recursive dwelling (the consideration and reconsideration of opportunity).

The performer gains access to the work through an elaborate Max/MSP 5 performance patch (see Figure 2) programmed by Reynolds's musical assistant for the project, Jacob David Sudol. This patch provides labels and controls for each "stage" in *AngelSpace*.

2. SPATIALIZATION

Both authors have extensively explored surround sound in nearly all their compositions that feature electronics. Following this interest, *AngelSpace* uses a six-speaker system symmetrically surrounds the audience (see Figure 2).

The *AngelSpace* patch directly routes each channel of the five six-channel computer "images" (D5 combined with D2, S8 combined with S7, and D10) to a corresponding channel in the six-speaker system. In contrast all the other sonic materials or audio files in *AngelSpace* have either one or two channels. The patch dynamically diffuses these latter materials in real-time following spatialization algorithms that the authors designed specifically for this composition.

[1] This essay contains text from the *AngelSpace* (2008-2009) technical score [4] that was written by both authors and additional text written by the first author.

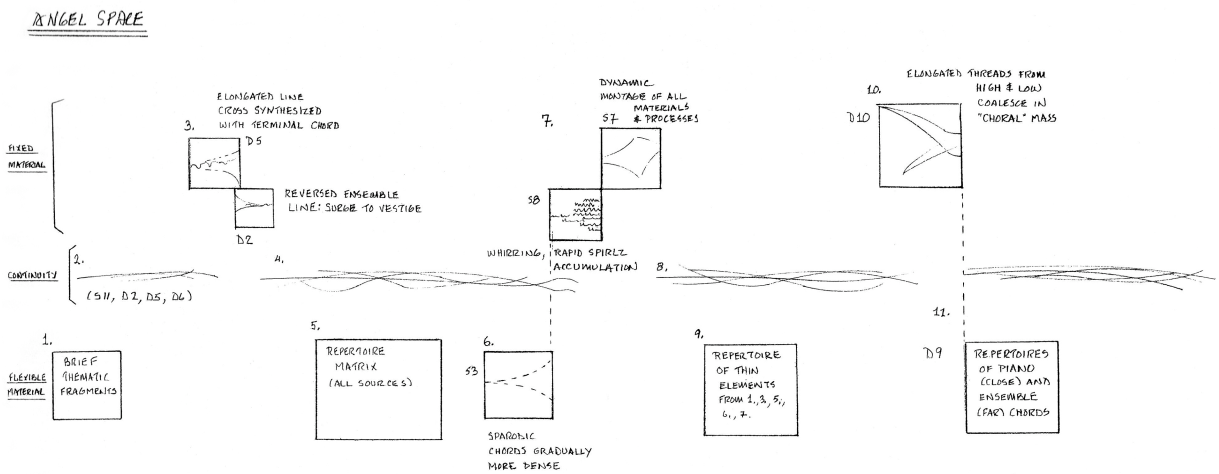

Figure 1. Overall plan for *AngelSpace*.

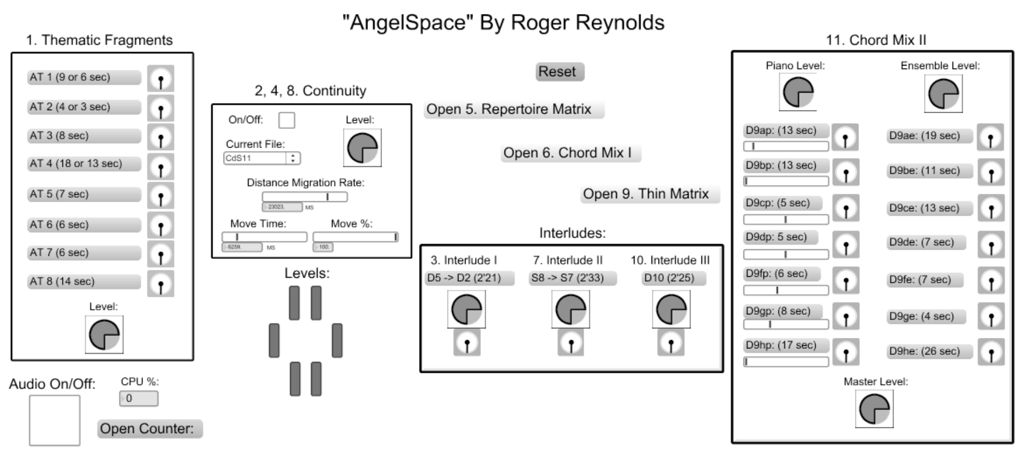

Figure 2. Screenshot of the Max/MSP 5 patch used to perform *AngelSpace*.

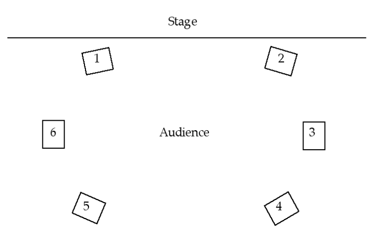

Figure 3. The six-speaker set up for *AngelSpace.*

2.1. Spatialization Algorithms

The spatial position for each one or two-channel audio file is normally defined by a six-speaker position within location set, a transition between two six-speaker positions, and a distance.

Location sets involve specific sequences of six-speaker positions (see Figure 2) as defined by groups "A" through "D" (see Table 1). When moving an audio begins at the first six-speaker position of randomly selected location set (e.g. speakers 4 and 5 in location set A).

Location Set	Sequence of six-speaker positions
A	45, 51, 63, 12, 632, 521, 416
B	51, 62, 13, 264, 153, 642, 531
C	62, 13, 24, 35, 4612, 3165, 2654
D	52, 63, 14, 25, 143, 521, 316

Table 1. Definition of location sets "A" through "D."

The audio then progresses through the sequence of six-speaker positions. At the sequence's end another location set is chosen at random. This process can then repeat indefinitely. A move time defines the time spent between six-speaker positions. A move time jitter defines a random variability of the move time. A move time percentage defines the amount of time that a move time is spent transitioning between two six-speaker positions as opposed to remaining at the initial position.

Distance is defined as a reverberant balance – e.g. a wet/dry ratio. The wet/dry ratio of 1 defines the greatest distance. In this case the reverberant sound of a one channel audio file is equally present in all speakers while that left and right channels of a stereo file are equally present in the respective three left (channels 1, 6, and 5) and three right (channels 2, 3, and 4) speakers. The absence of direct signal degrades the perception of position and its successions. A wet/dry ratio of 0 defines the greatest proximity. In this case only the direct signal from the location sets and transitions from one position to the next are heard.

In each section distance is either fixed for every sound file or dynamically determined by a distance migration rate. A distance migration rate is an automated cycle over which the apparent proximity of the sound file playing varies continuously between a predefined maximum proximity (e.g. a low wet/dry ratio) and maximum distance (e.g. a high wet/dry ratio). A distance migration rate can exist either in a single curve mode (see Figure 3) or a reciprocal curve mode (see Figure 4). If there is a single curve then the conditions established by the selected point alter continuously thereafter as described by the curve. However, if a pair of reciprocal curves is used, points are chosen randomly from either curve (as the reading moves through time). Associated events then remain fixed at the spatial distance location selected until completed.

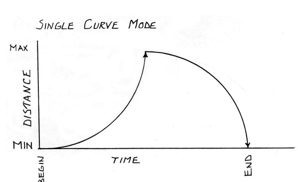

Figure 4. Graphic representation of the single curve mode.

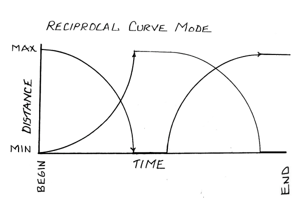

Figure 5. Graphic representation of the reciprocal curve mode.

In the initial implementation the distance migration rate curves moved linearly between the distance extremes. The authors found that this transition rate obscured one's perception of the transition in the single curve mode. After experimentation Sudol found that the distance migration cycle is far more perceptually present and satisfying when an exponent of 2 scales the transition from the maximum proximate position to the maximum distant position and when an exponent of .5 scales the transition from maximum distant position to the maximum proximate position.

In the initial implementation the pair of curves in the reciprocal curve mode moved between the maximum proximate and maximum distant positions without any pauses. This method gave the authors the impression that the distance position for each audio file is chosen completely at random. On the other hand a later inserted steady-state duration placed at each distance extreme gave the authors a much greater impression that each distance position was randomly chosen from a pair of reciprocal curves.

3. THE "STAGES" IN *ANGELSPACE*

3.1. Thematic Fragments

In the first "stage" of *AngelSpace* the performer introduces the desired fragments in leisurely isolation by clicking a corresponding rectangular box. These thematic fragments consist of eleven stereo audio files. The performer can directly choose to play five of the thematic fragments. The other six fragments comprise three pairs. When the performer chooses a pairs the computer randomly selects and plays one of the two files.

To ensure smooth transitions between each fragment the authors reverberated each audio file with the same stereo impulse response. The authors also carefully modified each file's gain to provide a satisfying balance among all the thematic fragments. During performance the first iteration of each fragment sounds at the original amplitude. Following this each successive iteration contains a 33% level reduction from the previous amplitude, so that a file's first iteration has an amplitude of 1., the second iteration an amplitude of .66, the third an amplitude of .33, and so on. The authors chose this specific level reduction to provide the impression that each specific thematic fragment gradually fades into the atmosphere or one's memory with each successive repetition.

All the thematic fragments sound at a relatively proximate distance. The performance patch randomly assigns each selected fragment one of four sets of move time, move time jitter, and move time percentage position controls. These sets control how the sound then spatially moves. Sudol created these four sets to introduce the listener to different ways that sound can move in *AngelSpace*. The authors specifically tuned the sets' values by ear.

During rehearsals of an approximately thirty minute performance of *AngelSpace* Reynolds advised Sudol to place significant silences between the presentation of each thematic fragment. Both authors found their dynamic engagement with the work's global form to largely improve when Sudol timed these silences by ear. Sudol's later analysis of these rehearsals' recordings suggests that the effect works best when the silences are between six to ten seconds long. During these same rehearsals Reynolds also advised Sudol to progress from the first to the last fragment with few repetitions of previously heard fragments. The authors found this advice sound after comparing performances where Sudol frequently repeated the fragments in one thirty- minute performance and rarely repeated the fragments in another. This said, since *AngelSpace* has no strictly prescribed duration, one could imagine that in a much longer duration that a performance with more repeated fragments would have a greater gestalt than a performance with less repeats. At this moment neither the authors nor any other performer has tested this hypothesis.

3.2. Continuity

Continuity provides a subtly evolving thread of sound that is always changing character and position. In the second, fourth, and eighth "stages" Continuity bridges between more assertive materials (e.g. the Thematic Fragments and Interlude I). It also sometimes remains active as a background while sparse and moderately assertive elements are superimposed on it (as with Thin Mix). This unique feature enters almost imperceptibly after the Thematic Fragments and then remains perceptually "there" until the last chord.

In Continuity there are three groups of quiet, mellifluous, and muted stereo audio files from *The Angel of Death* Themes. The *mellifluous* group comes from the S11 group and includes twelve samples ranging from 2456 to 12721 milliseconds. The *quiet* group comes from the D5 group and includes seven samples ranging 4000 to 12893 milliseconds. The *muted* group comes from the D6 group and includes nine samples that ranging 8000 to 27774. To ensure smooth transitions the authors only slightly modified the original amplitude of a few samples from within these groups.

Each sample mostly sounds in isolation; however as each file approaches its end the next sample smoothly enters. In each group the samples have a preferred order that in part determines what sample sounds next. Specifically each sample has its own specific probability table that the patch consults near a sample's end to determine which sample will follow. Figures 6 and 7 respectively show the probability tables for the first and twelfth samples in the *mellifluous* group. In both probability tables positions 0 through 11 on the x-axis correspond to the first through twelfth samples in the mellifluous group. Position 12 on the x-axis corresponds to a jump to either the *quiet* or *muted* group. The y-axis provides weightings for the position on the x-axis so that in Figure 6 there is no chance that the first sample will repeat, a large chance that the second sample will follow, and an equal probability that all the other samples or a jump to another group will follow. Each sample, besides the last, in a group has its own similar probability distributions (i.e. no chance of a repeated sample, high probability that the next sample will follow, etc.).

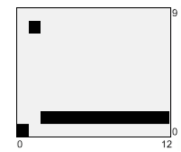

Figure 6. Probability table for the *mellifluous* group's first sample.

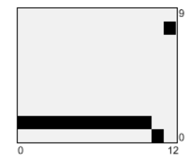

Figure 7. Probability table for the *mellifluous* group's twelfth sample.

For the last sample in a group a large chance that the next sample will come from another group replaces the possibility that the next sample within the group will follow – the metaphor being that a new group replaces the subsequent sample within a group at that group's end.

When the performer initiates Continuity the patch chooses a group at random. Similarly whenever a group is chosen the patch plays back a random a sample from that group. After extensively testing all the aforementioned continuity algorithms Reynolds found that Continuity occasionally spent an inordinately long duration in one group. To rectify this Sudol created and tuned a "bump time" function that automatically forces Continuity to jump to a new group after a specific time if the other algorithms have not yet caused a spontaneous jump to a new group. Each group has a preset "bump time" and, to ensure that the samples constantly overlap in the same manner, "bump time" does not cause the jump until the currently active sample approaches its end.

Distance is dynamically controlled by a distance migration rate in single curve mode. In the performance patch controls for the distance migrate rate's cycle duration, as well as for the move time and the move time percentage allow the performer to subtly adjust the nature of Continuity. In contrast a fixed move time jitter gives Continuity's spatialization a constant variability that the performer cannot modify. A rotary dial allows the performer to adjust the overall playback level of Continuity.

Once the Continuity process is initiated it can continue without performer input. Flexibility mainly primarily concerns whether it is or is not in operation and how loud it is. In rehearsals Sudol discovered that a slow fade in from silence in "stage" two benefits the desired aggregate impression of Continuity. While rehearsing a thirty-minute performance the authors have found that this fade in works well when it lasts between three to four minutes. In the other two sections when Continuity sounds in isolation (i.e. "stages" four and eight) Sudol has explored changing the spatialization parameters. The authors remain uncertain what effect these modifications have on the work's gestalt. In these same two "stages" the authors have not found that modifications of the Continuity's level to have any remarkable effect. However the authors have found that the performer must carefully adjust Continuity's level in other "stages" discussed below.

3.3. Interludes

The three six-channel Interludes ("stages" three, seven, and ten) have pre-established content and spatialization as well as a fixed duration. Reynolds constructed the first two interludes by overlapping and slightly modifying the levels of two computer "images" from *The Angel of Death*. For each Interlude the Max/MSP patch only allows the performer to initiate and modify the level of the corresponding sound file.

In rehearsals Reynolds commented that the three Interludes should provide a clear dramatic shape that contrasts the more open and stationary Thematic Fragments, Continuity, Repertoire Matrix, Thin Matrix, and Chord Mix II "stages." Responding to this comment Sudol discovered that one should initiate all three Interludes at their initial preset level as well as turn lower the Continuity's level by approximately 10 to 20% before initiating Interludes I and II. In addition, as all three Interludes progress one should progressively lower Continuity to a virtually inaudible level over the first quarter to third of the Interludes' duration as well as dynamically adjust each Interlude's level to exaggerate moments of calm stasis and climax. To facilitate this latter action Sudol normalized all the Interludes to about 90% so that one can raise the Interludes' levels to their maximum and not produce speaker distortion. Furthermore Sudol has also found careful study of each Interlude helps one forecast appropriate moments to make level adjustments for both Continuity and the Interludes. To insure smooth transitions out of the Interludes both authors have found it helpful to raise Continuity's level to approximately 75% over the last fifth of Interludes I and II and maintain a Continuity level of approximately 15% before the Interlude III ceases. In contrast Sudol found that it does not help to modify the Interludes' levels as they approach their ends. Finally Sudol has found it helpful to notice the materials that Continuity and the Interludes share. For example, in rehearsals and performance Sudol frequently brings in Interlude I soon after the Continuity presents the same material. This presentation provides a satisfyingly logical entrance point for Interlude I and reinforces that Continuity can move seamlessly alternate and transition between a perceptual foreground and background.

3.4. Chord Mix II

The last "stage" is an expansive progression of chords from Theme 5 in *The Angel of Death* (see Figure 8).

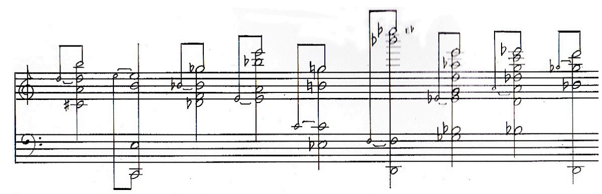

Figure 8. The chords associated with Chord Mix II.

The aim is to produce a dreamscape with stereo recordings of these chords as well as Continuity. The patch gives the performer access to two sets of these elongated, dream-like chords – unaltered Piano chords and Ensemble chords that are altered in pitch and time stretched. The duration of each chord is shown in its associated box. Each chord also has a clock dial to show the rate at which the file is being read to allow careful control of the resulting continuity between chords. The performer individually triggers each chord, nominally alternating between Piano and Ensemble chords as well as following the indicated ordering.

To ensure smooth transitions at between chords the authors reverberated each chord with the same stereo impulse response. Before this Sudol also added a fade out at the end of each Piano chord's resonance to create a smooth nearly unperceivable transition between the original files' dry resonances and the artificial reverb. The authors also modified a few files' gains to insure a satisfying balance among the chords. Similar to what happens in the Thematic Fragments, if a given Piano chord is repeated in performance, each instantiation will be reduced in level by 33%. In contrast, all instantiations of Ensemble chords sound at their original amplitudes.

Chord Mix II uses Villi Pulkki's VBAP algorithms and Max/MSP external [1] to create completely different positions than all the other "stages." Specifically the original stereo file chords and their two channels are spatialized so that their left and right channels appear at a 90-degree angle along the perimeter of the spatial field. The Piano chords move counter-clockwise and the Ensemble chords move clockwise. The speed of this rotation has a preset value that is directly proportional to the file's duration so that the longer the file the slower the rotation. The performer also has access to sliders that can dynamically alter the movement speed of each Piano chord.

Contrastingly Chord Mix II uses the method discussed in section 2.1 to create two fixed distances – a relatively proximate one for the Piano chords and a relatively distant one for the Ensemble chords. This latter method perceptually places the Piano chords in the spatial foreground and the Ensemble chords in the background.

During listening sessions the authors discovered that, even with the modifications mentioned above, each Piano chord's ending is too distractingly abrupt to be heard in isolation. In rehearsals the authors discovered that one can easily overcome this problem by initiating an Ensemble chord during the last fifth or sixth of a Piano chord's resonance. Through experiments the authors also discovered that Continuity should remain at a low 10 to 20% level during Chord Mix II. This level creates a clear hierarchical balance where the Piano chords exist in the foreground, the Ensemble chords in the background, and Continuity in the remote distance. Finally, during one of Sudol's rehearsals the authors also discovered that, contrary to Reynolds's original plan, *AngelSpace* ends most effectively when Continuity disappears about half-way through the last Ensemble chord. To foreshadow this disappearance Sudol typically incrementally turns down Continuity during the second half of Chord Mix II's duration.

3.5. Chord Mix I

The sixth "stage" consists of three groups of assertive chords (including both ensemble and piano versions) that the performer activates so that, over a total, prescribed acceleration time a very dense and intense, iterative texture accumulates. The performer enters the groups sequential and controls the force of their composite effect by determining the total acceleration time, or duration of the process, which layers are activated, and how many of the element sequences are allowed to continue repeating.

There are three register-identified groups of brief forceful chords derived from *The Angel of Death* Theme 2. The *high* group includes eight elements ranging from 73 to 638 milliseconds in duration. The *medium* group includes ten elements ranging from 364 to 1728 milliseconds. The *low* group includes eight elements ranging from 1057 to 5446 milliseconds. In each group successive elements become progressively longer and more forceful. When a group is selected each element enters sequentially following a predetermined order at nearly equal divisions of the remaining acceleration time. Each element, once introduced, repeats. The performer can cease each element's successive repetitions once it has been introduced; however, an element stopped thusly, cannot be been turned back on.

Each element has its own preordained repetition rate. In an early version, three master clocks, one for each group, controlled the repetition rate for each element in each group. These master clocks established a repetition rate range that progresses from the predetermined slowest rate when one instantiates a group to the fastest rate at the end of the acceleration time. After auditioning this approach Reynolds considered the effect too regular and predictable. Reynolds also found that this predictability particularly upsetting when only one group sounded at the beginning of the process. To rectify this Sudol gave each element its own repetition rate range. Sudol created these repetition rate ranges by progressively applying less of a random deviation to the aforementioned master clocks' ranges. The patch also contains a preset onset jitter that varies when each new element in a chosen group enters and a slider that allows the performer to dynamically apply a jitter on all the active repetition rates.

Because of the forceful and assertive character of Chord Mix I the authors did not alter the amplitudes of the original mono audio files. However, to make each successive repetition sound less mechanical and more realistic Sudol included an amplitude jitter that modifies each element's amplitude by a random decrement within a performed-controlled range.

There are two spatialization modes for Chord Mix I. In both modes each audio file sounds at a fixed position and distance. In the first mode each successive element progresses sequentially through the location sets in the method described in section 2.1 and queries its distance from a distance migration cycle in the reciprocal curve mode. In the second mode *high* elements inhabit the front-left third of the spatial field, *medium* elements inhabit the front-right third of the spatial field, and *low* elements inhabit the back third of the spatial field. For the second mode a dry/wet ratio of 1.0 places all elements at the maximum possible proximity. Positions in the second spatialization mode come from the six-speaker positions shown in Table 1. For this mode Sudol also eliminated positions that exist equally in two or three thirds of the spatial field. At the beginning of Chord Mix I there is a 100% probability the first mode determines the spatialization. At the end of the acceleration time there is a 100% probability that the second mode determines the spatialization. The probability linearly moves between these two extremes over the course of the acceleration time.

In Chord Mix I the performer needs to carefully regulate and shape a dramatic dynamic and iterative swell. During experiments Sudol discovered that this process works best when the acceleration time is two to three minutes long. Both authors found that shorter durations progress too quickly. Reynolds suggested that the Chord Mix I would work best when one starts with the *middle* group, instantiates the *high* group approximately one-third of the way into the acceleration time, and then instantiates the *low* group approximately two-thirds of the way into the acceleration time. Both authors found that the process works best when one to four elements in any group repeat. In addition Reynolds decided that one should turn off elements in the order that they enter. During rehearsals Sudol has found that the entire process works best when two to three elements in any group repeat in the beginning and four or possibly even five elements repeat at the end of the process. Sudol also found that one can further prolong the Chord Mix I by progressively increasing the amplitude and repetition jitters after the acceleration time has completed.

In addition the performer must seamlessly transition from Chord Mix I to Interlude II. Sudol discovered that careful control of Continuity's level aids this transition. Specifically one should Continuity keep at approximately 75% for the first third of Chord Mix I, reduce it to approximately 50% during the second third, and then swiftly raise it to 100% just before ceasing Chord Mix I. Soon after Chord Mix I ends the performer should instantiate Interlude II and significantly turn down Continuity. Sudol also discovered that this transition works best if it is the only time Continuity sounds at its maximum level.

3.6. Repertoire Matrix

This texture is defined by the selection of groups of elements that have themselves a characteristic nature (see Figure 9).

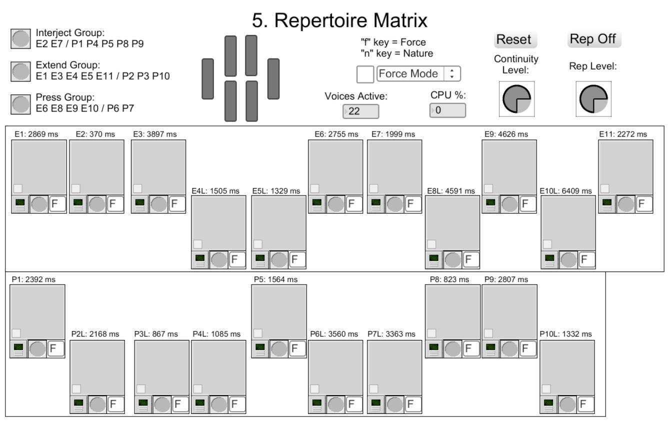

Figure 9. Screenshot of the Repertoire Matrix.

These elements come from a mix of piano and ensemble versions of similar musical ideas. Elements from a selected group enter according to pre-established conditions and repeat continuously with some omissions. Once a satisfying pattern of repeating elements is established, the performer can alter the spectral content of the various components. Adjustment of textural and timbral character can continue in alternation with changes in the basic repetitive texture, as can changes of the group from which elements are drawn.

There are groups each containing a unique selection of piano and ensemble stereo musical fragments from *The Angel of Death* Themes. The *interject* group includes seven elements ranging from 400 to 2800 milliseconds. The *extend* group includes eight elements ranging from 900 to 3900 milliseconds. The *push* group includes six elements ranging from 750 to 6400 milliseconds. A graphic layout associates a square field with each group element. The upper rank of each stratum (in either the Ensemble or Piano fields) has more assertive elements, the lower, less assertive elements. To aid with balancing Sudol originally attempted to give all the elements equal amplitudes. Unfortunately some of the original audio files had rather poor signal-to-noise ratios. Rather than delete these files the authors decided to assign them a less assertive role.

The defined dimensions of each element's square field can be switched between two modalities – force and nature. During performance the performer can modify an element's character by adjusting the positions of a square cursor in each square field. The performer, however, cannot dynamically change an element's dimensional attributes while that specific element sounds.

For the default force mode the y-axis defines the element's amplitude level and the x-axis defines the probability of repetition (see Figure 10).

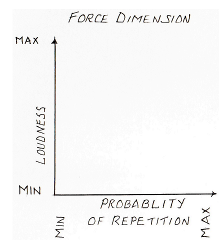

Figure 10. Definition of a square field's force dimension.

A minimum level effectively turns off an element until the performer raises the level. Each group has a different range of values for the x-axis. Besides these user-defined controls global preset repetition and amplitude jitters add a lifelike amount of irregularity to each element's repetitions. In addition an omission rate will omit between one to three repetitions after each file has repeated seven to nine times or four to nine times if an element's repetitions have not yet been omitted. The author's tuned these omissions so that the listener perceives each element's repetitions creating a clearly segmented musical phrase.

In the nature mode a brick-wall FFT filter notches a triangle out each file's spectral map, creating two equal and complimentary high and low components. This triangle filtering is applied only from 0 to 3 kHz. All nature modifications only apply to the low component. For the nature mode, the y-axis defines a frequency shift with a range from 0 to -500 Hz and the x-axis defines a temporal delay of the lower component in relationship to the upper with a range from 0 to 10% of each element's total duration. By default the nature mode performs no modifications; therefore, only the performer can modify the spectral content of each element.

When the performer selects a group a set of predetermined entrance and repetition times are applied to a random ordering of the group's elements. Each element enters with its own predetermined amplitude. The authors intuitively composed these presets for each group and element by ear. Once an element has begun the performer can alter its dimensional attributes at will. When the performer selects a different group the first element of the previous group will become inactive. When the second element of the new group enters the second element of the previous group will stop repeating. This process continues until either all the previous group's elements are inactive or all the new group's elements have entered. In the latter case all remaining elements from the previous group are inactivated.

Repertoire Matrix and Chord Mix I use a nearly identical spatialization approach to determine each file's distance and position. The difference is that each piano sounds' positions come from the location sets A and B while each ensemble sounds' come from the location sets C and D.

The Repertoire Matrix provides the performer with more controls or flexibility than any other "stage." Fortunately the authors have found that this flexibility usually produces interesting results. The authors have only made a few observations about successful ways to progress. For example, Sudol has found it helps to reduce Continuity's level to approximately 50% for most the duration of the Repertoire Matrix. Both authors have enjoyed rehearsed performances when Sudol allowed each element in a group to enter before he slowly altered each element's nature as much as possible. Sudol has also found it effective to end the Repertoire Matrix by progressively turning off each active element.

3.7. Thin Matrix

The Thin Matrix includes three groups of relatively continuous instrumental phrases. These sound files have been modified by continuously changing pair of band-reject brick-wall FFT filters that simulate formants in a metaphorically negative way where "emphasis" equals absence. This approach thins the spectral energy and nature of each in constantly varying way. The aim is to create an amalgam of familiar materials whose character has been strongly altered so that they become pale and transparent. This allows Continuity to be equally present for the entire duration of Thin Matrix.

There are three groups of four to six stereo source files. The files range in duration from 9858 to 16042 milliseconds. Reynolds created a series of probabilistic rules that determined how the filters modified the original files. Sudol programmed these and the authors rigorously tried out different parameters for these rules before eventually choosing fifteen thinned files. Finally Sudol reverberated each audio file with the same stereo impulse response to create smooth transitions between the elements and Continuity.

Thin Matrix and Repertoire Matrix operate and progress in nearly identical manners. One distinction is that Thin Matrix only has the force mode. In addition, amplitude can be dynamically controlled during an element's playback in Thin Matrix. The other distinction is that Thin Matrix elements move in a predetermined manner that is nearly identical to the spatialization method used in Continuity – the differences being Thin Matrix elements move quicker and have fixed distances that are determined by a distance migration cycle in the reciprocal curve mode.

The Thin Matrix allows the performer great flexibility in performance. Fortunately the authors have discovered that this almost always leads to intriguing musical results when Thin Matrix and Continuity have similar amplitudes.

4. CONCLUSION

To the authors' knowledge there are few compositions that dynamically and probabilistically explore the meaning and implication of complex materials in a manner similar to *AngelSpace*. Furthermore, the authors do not know of any other works that explore the spectral modification techniques used in the Repertoire and Thin Matrixes. Since both these formal and digital signal processing explorations proved so fruitful for this project the authors' predict these approaches and will prove to be robust and fruitful to the authors and others in the future. Furthermore it is the authors' hope that *AngelSpace* and similar computer music compositions featuring performer-controlled probabilistic algorithms will provide composers, performers, and listeners alike with new and exciting ways to understand and appreciate musical processes and materials.

5. REFERENCES

[1] Pulkki, V. "Compensating displacement of amplitude-panned virtual sources," in *Proceedings of the Audio Engineering Society 22th Internation Conference on Virtual, Synthetic and Entertainment Audio*, Espoo, Finland 2002, pp. 186-195.

[2] Reynolds, Roger. *The Angel of Death*. New York: Editions Peters, 2001.

[3] Reynolds, Roger. *The Image Machine*. New York: Editions Peters, 2005.

[4] Reynolds, Roger & Jacob Sudol. *AngelSpace: An introduction to the concepts and performance practice bases of a real-time performance composition.* New York: Editions Peters, October 2009.